
Budgeted Online Influence Maximization

Pierre Perrault^{1 2 3} Jennifer Healey¹ Zheng Wen⁴ Michal Valko^{4 3}

Abstract

We introduce a new budgeted framework for online influence maximization, considering the total cost of an advertising campaign instead of the common cardinality constraint on a chosen influencer set. Our approach models better the real-world setting where the cost of influencers varies and advertisers want to find the best value for their overall social advertising budget. We propose an algorithm assuming an independent cascade diffusion model and edge-level semi-bandit feedback, and provide both theoretical and experimental results. Our analysis is also valid for the cardinality-constraint setting and improves the state of the art regret bound in this case.

1. Introduction

Viral marketing through online social networks now represents a significant part of many digital advertising budgets. In this form of marketing, companies incentivize chosen influencers in social networks (e.g., Facebook, Twitter, YouTube) to feature a product in hopes that their followers will adopt the product and repost the recommendation to their own network of followers. The effectiveness of the chosen set of influencers can be measured by the expected number of users that adopt the product due to their initial recommendation, called the *spread*. *Influence maximization* (IM, Kempe et al., 2003) is the problem of choosing the optimal set of influencers to maximize the spread under a cardinality constraint on the chosen set.

In order to define the spread, we need to specify a diffusion process such as independent cascade (IC) or linear threshold (LT) (Kempe et al., 2003). The parameters of these models are usually *unknown*. Different methods exist to estimate the parameters of the diffusion model from historical data (see section 1.1) however historical data is often difficult to obtain. Another possibility is to consider *online influ-*

ence maximization (OIM) (Vaswani et al., 2015; Wen et al., 2017) where an agent actively learns about the network by interacting with it repeatedly, trying to find the best seed influencers. The agent thus faces the dilemma of exploration versus exploitation, allowing us to see it as *multi-armed bandits* problem (Auer et al., 2002). More precisely, the agent faces IM over T rounds. Each round, it selects m seeds (based on *feedback* from prior rounds) and diffusion occurs; then it gains a reward equal to the spread and receives some feedback on the diffusion.

IM and OIM optimize with the constraint of a fixed number of seeds. This reflects a fixed seed cost model, for example, where influencers are incentivized by being given an identical free product. In reality, however, many influencers demand different levels of compensation. Those with a high out-degree (e.g., number of followers) are usually more expensive. Due to these cost variations, marketers usually wish to optimize their seed sets S under a budget $c(S) \leq b$ rather than a cardinality constraint $|S| \leq m$. Optimizing a seed set under a budget has been studied in the *offline* case by Nguyen and Zheng (2013). In the *online* case, Wang et al. (2020) considered the relaxed constraint $\mathbb{E}[c(S)] \leq b$, where the expectation is over the possible randomness of S .¹ We believe however that the constraint of a fixed, equal budget $c(S) \leq b$ at each round does not sufficiently model the willingness to choose a cost-efficient seed set. Indeed, we see that the choice of b is crucial: a b too large translates into a waste of budget (some seeds that are too expensive will be chosen) and a b too small translates into a waste of time (a whole round is used to influence only a few users). To circumvent this issue, instead of a budget per round, in our framework, we allow the agent to choose seed sets of any cost at each round, under an overall budget constraint (equal to $B = bT$ for instance). In summary, we incorporate the OIM framework into a *budgeted bandit* setting. Our setting is more flexible for the agent, and better meets real-world needs.

1.1. Related work on IM

IM can be formally defined as follows. A social network is modeled as a directed graph $G = (V, E)$, with nodes V

¹This relaxation is to avoid a computationally costly *partial enumeration* (Krause and Guestrin, 2005; Khuller et al., 1999)

¹Adobe Research, San Jose, CA ²ENS Paris-Saclay ³Inria Lille ⁴DeepMind. Correspondence to: Pierre Perrault <pierre.perrault@outlook.com>.

representing users and edges E representing connections. An underlying diffusion model D governs how information spreads in G . More precisely, D is a probability distribution on subgraphs G' of G , and given some seed set S , the spread $\sigma(S)$ is defined as the expected number of S -reachable² nodes in $G' \sim D$. IM aims to find S that is a solution to

$$\max_{|S|=m} \sigma(S). \quad (1)$$

Although IM is NP-hard under standard diffusion models — i.e., IC and LT — σ is a monotone *submodular*³ function (Fujishige, 2005), and given a value oracle access to σ , the standard GREEDY algorithm solves (1) within a $1 - 1/e$ approximation factor (Nemhauser et al., 1978). There have been multiple lines of work for IM, including the development of heuristics, approximation algorithms, as well as alternative diffusion models (Leskovec et al., 2007; Goyal et al., 2011; Tang et al., 2014; 2015). Additionally, there are also results on learning D from data in the case it is not known (Saito et al., 2008; Goyal et al., 2010; Gomez-Rodriguez et al., 2012; Netrapalli and Sanghavi, 2012).

1.2. Related work on OIM

Prior work in OIM has mainly considered either *node level semi-bandit* feedback (Vaswani et al., 2015), where the agent observes all the S -reachable nodes in G' , or *edge level semi-bandit* feedback (Wen et al., 2017), where the agent observes the whole S -reachable subgraph (i.e., the subgraph of G' induced by S -reachable nodes). Other, weaker, feedback settings have also been studied including: pairwise influence feedback, where all nodes that would be influenced by a seed set are observed but not the edges connecting them, i.e., $(\{i\}$ -reachable nodes) $_{i \in S}$ is observed (Vaswani et al., 2017); local feedback, where the agent observes a set of out-neighbors of S (Carpentier and Valko, 2016) and immediate neighbor observation where the agent only observes the out-degree of S (Lugosi et al., 2019).

1.3. Our contributions

In this paper, we define the budgeted OIM paradigm and propose a performance metric for an online policy on this problem using the notion of *approximation regret* (Chen et al., 2013). To the best of our knowledge, the both of contributions are new. We then focus our study on the IC model with edge level semi-bandit feedback. We design a CUCB-style algorithm and prove logarithmic regret bounds. We also propose some modifications of this algorithm with improving the regret rates. These gains apply to the non-budgeted setting, giving an improvement over the state-of-the-art analysis of the standard CUCB-approach (Wang and Chen, 2017). Our proof incorporates an approximation

guarantee of GREEDY for ratio of submodular and modular functions, which may also be of independent interest.

2. Problem definition

In this section, we formulate the problem of budgeted OIM and give a regret definition for evaluating policies in that setting. We also justify our choice for this notion of regret. We typeset vectors in bold and indicate components with indices, e.g., for some set I , $\mathbf{a} = (a_i)_{i \in I} \in \mathbb{R}^I$ is a vector on I . Let \mathbf{e}_i be the i^{th} canonical unit vector of \mathbb{R}^I . The incidence vector of any subset $A \subset I$ is $\mathbf{e}_A \triangleq \sum_{i \in A} \mathbf{e}_i$.

We consider a fixed directed network $G = (V, E)$, known to the agent, with $V \triangleq \{1, \dots, |V|\}$. We denote by $ij \in E$ the directed edge from node i to j in G . We assume that G doesn't have self-loops, i.e., for all $ij \in E$, $i \neq j$. For a node $i \in V$, a subset $S \subset V$, and a vector $\mathbf{w} \in \{0, 1\}^E$, the predicate $S \overset{\mathbf{w}}{\rightsquigarrow} i$ holds if, in the graph defined by $G_{\mathbf{w}} \triangleq (V, \{ij \in E, w_{ij} = 1\})$, there is a forward path from a node in S to the node i . If it holds, we say that i is influenced by S under \mathbf{w} . We define $p_i(S; \mathbf{w}) \triangleq \mathbb{I}\{S \overset{\mathbf{w}}{\rightsquigarrow} i\}$ and the *spread* as $\sigma(S; \mathbf{w}) \triangleq \left| \left\{ i \in V, S \overset{\mathbf{w}}{\rightsquigarrow} i \right\} \right|$. Our diffusion process is defined by the random vector $\mathbf{W} \in \{0, 1\}^E$, and our cost is defined by the random⁴ vector $\mathbf{C} \in [0, 1]^{V \cup \{0\}}$ where the added component C_0 represents any fixed costs⁵. Notice, random costs are neither assumed to be mutually independent nor independent from \mathbf{W} . We will see that components of \mathbf{W} might however be mutually independent (e.g., for the IC model).

2.1. Budgeted online influence maximization

The agent interacts with the diffusion process across several rounds, using a learning policy. At each round $t \geq 1$, the agent first selects a seed set $S_t \subset V$, based on its past observations. Then, the random vectors for both the diffusion process $\mathbf{W}_t \sim \mathbb{P}_{\mathbf{W}}$ and the costs $\mathbf{C}_t \sim \mathbb{P}_{\mathbf{C}}$ are sampled independently from previous rounds. Then, the agent observes some feedback from both the diffusion process and the costs.

We provide in (2) the expected cumulative rewards F_B defined for some total budget $B > 0$. The goal for the agent is to follow a learning policy π maximizing F_B . In (2), recall that S_t is the seed set selected by π at round t .

⁴Although costs are usually deterministic, we assume randomness for more generality (influencer campaigns may have uncertain surcharges for example).

⁵We provide a toy example where C_0 models a concrete quantity: Assume you want to fill your restaurant. You may pay some seeds and ask them to advertise/influence people. C_0 represents the cost of the food, the staff, the rent, the taxes, ...

²nodes that are reachable from some node in S .

³ f is submodular if $f(A \cup \{i\}) - f(A)$ is non-increasing with A .

$$F_B(\pi) \triangleq \mathbb{E} \left[\sum_{t=1}^{\tau_B-1} \sigma(S_t; \mathbf{W}_t) \right]. \quad (2)$$

τ_B is the random round at which the remaining budget becomes negative: if $B_t \triangleq B - \sum_{t' \leq t} (\mathbf{e}_{S_{t'}}^\top \mathbf{C}_{t'} + C_{0,t'})$, then $B_{\tau_B-1} \geq 0$ and $B_{\tau_B} < 0$. Notice, quantities B_t and τ_B are usual in budgeted multi-armed bandits (Xia et al., 2016; Ding et al., 2013).

2.2. Performance metric

We restrict ourselves to *efficient* policies, i.e., we consider a complexity constraint on the policy the agent can follow: For a round t , the space and time complexity for computing S_t has to be polynomial in $|V|$, and polylogarithmic in t . To evaluate the performance of a learning policy π , we use the notion of approximation regret (Kakade et al., 2009; Streeter and Golovin, 2009; Chen et al., 2016). The agent wants to follow a learning policy π which minimizes

$$R_{B,\varepsilon}(\pi) \triangleq (1 - 1/e - \varepsilon)F_B^* - F_B(\pi),$$

where F_B^* is the best possible value of F_B over all policies (thus leveraging on the knowledge of $\mathbb{P}_{\mathbf{W}}$ and $\mathbb{P}_{\mathbf{C}}$), and where $\varepsilon > 0$ is some parameter the agent can control to determine the tradeoff between computation and accuracy.

Remark 1. *This OIM with a total budget B is different from OIM in previous work, such as Wang and Chen (2017), even when we set all costs to be equal. In our setting, there is only one total budget for all rounds, and the policy is free to choose seed sets of different cost in each round, whereas in the previous work, each round had a fixed budget for the number/cost of seeds selected. Our setting thus avoid the use of a budget per round, which is in practice more difficult to establish than a global budget B . Nevertheless, as we will see in section 6, both types of constraints (global and per round) can be considered simultaneously when the true costs are known.*

2.3. Justification for the approximation regret

In the non-budgeted OIM problem with a cardinality constraint given by $m \in [|V|]$, let us recall that the approximation regret

$$R_{T,\varepsilon}(\pi) \triangleq \sum_{t \leq T} \max_{\substack{S \subset V, \\ |S|=m}} \mathbb{E}[(1 - 1/e - \varepsilon)\sigma(S; \mathbf{W}) - \sigma(S_t; \mathbf{W})]$$

is standard (Wen et al., 2017; Wang and Chen, 2017). In this notion of regret, the factor $(1 - 1/e - \varepsilon)$ (Feige, 1998; Chen et al., 2010) reflects the difficulty of approximating the following NP-Hard (Kempe et al., 2003) problem in the case the distribution $\mathbb{P}_{\mathbf{W}}$ is described by IC or LT, and is known to the agent:

$$\max_{S \subset V, |S|=m} \mathbb{E}[\sigma(S; \mathbf{W})]. \quad (3)$$

For our budgeted setting, at first sight, it is not straightforward to know which approximation factor to choose. Indeed, since the random horizon may be different in $F_B(\pi)$ and in F_B^* , the expected regret $R_{B,\varepsilon}(\pi)$ is not expressed as the expectation of a sum of approximation gaps, so we can't directly reduce the regret level approximability to the gap level approximability. We thus consider a quantity provably close to $R_{B,\varepsilon}(\pi)$ and easier to handle.

Proposition 1. *Define*

$$\lambda^* \triangleq \max_{S \subset V} \frac{\mathbb{E}[\sigma(S; \mathbf{W})]}{\mathbb{E}[\mathbf{e}_{S \cup \{0\}}^\top \mathbf{C}]}$$

For all $S \subset V$, define the gap corresponding to S as

$$\Delta(S) \triangleq (1 - 1/e - \varepsilon)\lambda^* \mathbb{E}[\mathbf{e}_{S \cup \{0\}}^\top \mathbf{C}] - \mathbb{E}[\sigma(S; \mathbf{W})].$$

Then, for any policy π selecting S_t at round t ,

$$\left| R_{B,\varepsilon}(\pi) - \mathbb{E} \left[\sum_{t=1}^{\tau_B-1} \Delta(S_t) \right] \right| \leq 2|V| + 2\lambda^*(1 + |V|).$$

From Proposition 1, whose proof can be found in Appendix A, $R_{B,\varepsilon}(\pi)$ and $\mathbb{E} \left[\sum_{t=1}^{\tau_B-1} \Delta(S_t) \right]$ are equivalent in term of regret upper bound rate. Therefore, the factor $(1 - 1/e - \varepsilon)$ should reflect the approximability of

$$\max_{S \subset V} \mathbb{E}[\sigma(S; \mathbf{W})] / \mathbb{E}[\mathbf{e}_{S \cup \{0\}}^\top \mathbf{C}] = \max_{S \subset V} f(S)/c(S). \quad (4)$$

Considering the specific problem where the cost function is of the form $c(S) = c_1|S| + c_0$, for some $(c_0, c_1) \in [0, 1]^2$, we can reduce the approximability of (4) to the approximability of the following problem considered in Wang et al. (2020):

$$\max_{S \subset V} \mathbb{E}[f(S)] \text{ s.t. } \mathbb{E}[|S|] \leq m, \quad (5)$$

for some given integer m , where the expectations are with respect to a randomization in the approximation algorithm. Wang et al. (2020) proved that this problem is NP-hard by reducing to the *set cover* problem. We show here that an approximation ratio α better than $1 - 1/e$ yields an approximation for set cover within $(1 - \delta) \log(|V|)$, $\delta > 0$, which is impossible unless $\text{NP} \subset \text{BPTIME}(n^{\mathcal{O}(\log \log(|V|))})$ (Feige, 1998). Consider the graph where the collection of out-neighborhoods is exactly the collection of sets in the set cover instance. First, trying out all possible values of m , we concentrate on the case in which the optimal m for set cover is tried out. As in Feige (1998), for $k \in \mathbb{N}^*$, we repeatedly apply the algorithm that α -approximate (5). It outputs a set S_k (that can be associated with a set of neighborhoods) and after each application the nodes already covered by previous applications are removed from the graph, giving a sequence of objective functions (f_k) with $f_1 = f$. We thus obtain

$$\mathbb{E}[f_k(S_k) | S_1, \dots, S_{k-1}] \geq \alpha \left(|V| - \sum_{k'=1}^{k-1} f_{k'}(S_{k'}) \right).$$

Noticing that $\mathbb{E}[f(S_1 \cup \dots \cup S_k)] = \sum_{k'=1}^k \mathbb{E}[f_{k'}(S_{k'})]$, we get

$$\mathbb{E}[f(S_1 \cup \dots \cup S_k)] \geq (1 - (1 - \alpha)^k)|V|.$$

After $\ell = \lceil \log(1/|V|)/\log(1 - \alpha) \rceil < (1 - \delta) \log(|V|)$ iterations, we obtain that $S = S_1 \cup \dots \cup S_\ell$ is a cover, i.e., $f(S) = |V|$. The result follows noticing that in expectation (and so with positive probability), we have $|S| \leq \ell m$.

3. Algorithm for IC with edge level semi-bandit feedback

3.1. Setting

For $\mathbf{w} \in [0, 1]^V$, we recall that we can define an IC model by taking $\mathbb{P}_{\mathbf{w}} = \otimes_{ij \in E} \text{Bernoulli}(w_{ij})$. We can extend the two previous functions p_i and σ to \mathbf{w} taking values in $[0, 1]^V$ as follows: Let $\mathbf{W} \sim \otimes_{ij \in E} \text{Bernoulli}(w_{ij})$. We define the probability that i is influenced by S under \mathbf{W} as $p_i(S; \mathbf{w}) \triangleq \mathbb{P}\left[S \stackrel{\mathbf{W}}{\rightsquigarrow} i\right]$, and we let the spread be $\sigma(S; \mathbf{w}) \triangleq \mathbb{E}\left[\left|\left\{i \in V, S \stackrel{\mathbf{W}}{\rightsquigarrow} i\right\}\right|\right]$. Another expression for the spread is $\sigma(S; \mathbf{w}) = \sum_{i \in V} p_i(S; \mathbf{w})$. We fix a weight vector on E , $\mathbf{w}^* \triangleq (w_{ij}^*)_{ij \in E} \in [0, 1]^E$, a cost vector on $V \cup \{0\}$, $\mathbf{c}^* \triangleq (c_i^*)_{i \in V \cup \{0\}} \in [0, 1]^{V \cup \{0\}}$, with $c_0^* > 0$. These quantities are initially *unknown* to the agent. We assume from now that

$$\mathbb{P}_{\mathbf{w}} \triangleq \otimes_{ij \in E} \text{Bernoulli}(w_{ij}^*),$$

and that

$$\mathbb{E}[\mathbf{C}] = \mathbf{c}^*.$$

We also define $S^* \in \arg \max_{S \subseteq V} \sigma(S; \mathbf{w}^*)/e_{S \cup \{0\}}^\top \mathbf{c}^*$. We assume that the feedback received by the agent at round t is $\left\{W_{ij,t}, ij \in E, S_t \stackrel{\mathbf{W}^t}{\rightsquigarrow} i\right\}$. The agent also receives semi-bandit feedback from the costs, i.e., $\{C_{i,t}, i \in V, i \in S_t \cup \{0\}\}$ is observed.

3.2. Algorithm design

In this subsection, we present BOIM-CUCB, CUCB for Budgeted OIM problem as Algorithm 1. As we saw in Proposition 1, the policy that, at each round, $(1 - 1/e - \varepsilon)$ -approximately maximize

$$S \mapsto \frac{\sigma(S; \mathbf{w}^*)}{e_S^\top \mathbf{c}^* + c_0^*} \quad (6)$$

has a bounded regret. Thus, BOIM-CUCB shall be based on this objective. Not only there are some estimation concerns due to the unknown parameters \mathbf{w}^* , \mathbf{c}^* , but in addition to that, we also need to evaluate/optimize our estimates of (6).

We begin by introducing some notations. We define the empirical means for $t \geq 1$ as: For all $i \in V \cup \{0\}$,

$$\bar{c}_{i,t-1} \triangleq \frac{\sum_{t' \in [t-1]} \mathbb{I}\{i \in S_{t'} \cup \{0\}\} C_{i,t'}}{N_{\oplus i, t-1}},$$

Algorithm 1 BOIM-CUCB

Input: $\varepsilon > 0, B_0 = B > 0$.

for each round $t \geq 1$ **do**

If true costs are known, then $\mathbf{c}_t \leftarrow \mathbf{c}^*$.

Compute S_t given by Algorithm 2 with input $S \mapsto \sigma(S; \mathbf{w}_t), \mathbf{c}_t$.

Select seed set S_t , and pay $e_{S_t \cup \{0\}}^\top \mathbf{C}_t$ (i.e., remove this cost from B_{t-1} to get the new budget B_t).

if $B_t \geq 0$, **then**

Get the reward $\sigma(S_t; \mathbf{W}_t)$, get the feedback, and update corresponding quantities accordingly.

else

The budget is exhausted: leave the for loop.

end if

end for

and for all $ij \in E$,

$$\bar{w}_{ij,t-1} \triangleq \frac{\sum_{t' \in [t-1]} \mathbb{I}\left\{S_{t'} \stackrel{\mathbf{W}^{t'}}{\rightsquigarrow} i\right\} W_{ij,t'}}{N_{\oplus i, t-1}},$$

where $N_{\oplus i, t-1} \triangleq \sum_{t'=1}^{t-1} \mathbb{I}\{i \in S_{t'}\}$, $N_{\oplus i, t-1} \triangleq \sum_{t'=1}^{t-1} \mathbb{I}\left\{S_{t'} \stackrel{\mathbf{W}^{t'}}{\rightsquigarrow} i\right\}$. Using concentration inequalities, we get confidence intervals for the above estimates. We are then able to use an upper-confidence-bound (UCB) strategy (Auer et al., 2002). More precisely, in the case costs are unknown, we first build the *lower* confidence bound (LCB) on c_i^* as follows

$$c_{i,t} \triangleq 0 \vee \left(\bar{c}_{i,t-1} - \sqrt{\frac{1.5 \log(t)}{N_{\oplus i, t-1}}} \right).$$

We can also define UCBs for w_{ij}^* :

$$w_{ij,t} \triangleq 1 \wedge \left(\bar{w}_{ij,t-1} + \sqrt{\frac{1.5 \log(t)}{N_{\oplus i, t-1}}} \right).$$

We use $w_{ij,t} = 1$ (and $c_{i,t} = 0$) when the corresponding counter is equal to 0. Our BOIM-CUCB approach chooses at each round t the seed set S_t given by Algorithm 2 which, as we shall see, approximately maximize $S \mapsto \sigma(S; \mathbf{w}_t)/\left(e_{S \cup \{0\}}^\top \mathbf{c}_t\right)$. Indeed, with high probability, this set function is an upper bound on the true ratio (6) (using that σ is non decreasing w.r.t. \mathbf{w}). Notice that this approach is followed by Wang and Chen (2017) for the non budgeted setting, i.e., they choose $S_t, |S_t| \leq m$ that approximately maximize $S \mapsto \sigma(S; \mathbf{w}_t)$. To complete the description of our algorithm, we need to describe Algorithm 2. This is the purpose of the following.

3.3. Greedy for ratio maximization

In BOIM-CUCB, one has to approximately maximize the ratio $S \mapsto \sigma(S; \mathbf{w}_t)/e_{S \cup \{0\}}^\top \mathbf{c}_t$, that is a ratio of submodular over modular function. A GREEDY technique can be

Algorithm 2 GREEDY for ratio, Lazy implementation

Input: σ that is an increasing submodular function,
 $\mathbf{c} \in [0, 1]^{V \cup \{0\}}$.
 $S_0 \leftarrow \emptyset$.
 $\rho \leftarrow [(\infty, i)]_{i \in V}$.
for $k \in [|V|]$ **do**
 checked $\leftarrow S_{k-1}$.
 (*) Remove the first element $\rho[0] = (\sim, i)$ from ρ .
 if $i \notin$ checked **then**
 Insert $((\sigma(\{i\} \cup S_{k-1}) - \sigma(S_{k-1}))/c_i, i)$ in ρ ,
 such that $\rho[0]$ is sorted in decreasing order.
 Add i to checked and go back to (*).
 else
 $S_k \leftarrow S_{k-1} \cup \{i\}$.
 end if
end for
 $k' \leftarrow \arg \max_{k \in \{0, \dots, |V|\}} \sigma(S_k) / \mathbf{e}_{S_k \cup \{0\}}^\top \mathbf{c}$.
Output: $S_{k'}$.

used (see Algorithm 2). Indeed, instead of maximizing the marginal contribution at each time step, as the standard GREEDY algorithm do, the approach is to maximize the so called bang-per-buck, i.e., the marginal contribution divided by the marginal cost. This builds a sequence of increasing subsets, and the final output is the one that maximizes the ratio. We prove in Appendix D the following Proposition 2, giving an approximation factor of $1 - 1/e$ for Algorithm 2.

Proposition 2. *Algorithm 2 with input σ, \mathbf{c} is guaranteed to obtain a solution S such that:*

$$(1 - e^{-1}) \frac{\sigma(S^*)}{\mathbf{e}_{S^* \cup \{0\}}^\top \mathbf{c}} \leq \frac{\sigma(S)}{\mathbf{e}_{S \cup \{0\}}^\top \mathbf{c}}.$$

Notice, a similar result as Proposition 2 is stated in Theorem 3.2 of Bai et al. (2016). However, their proof doesn't hold in our case, since their inequality (16) would be true only for a *normalized* cost (i.e. $c_0 = 0$). Actually, $c_0 = 0$ implies that S^* is a singleton, from subadditivity of σ .

For more efficiency, we use a greedy algorithm with lazy evaluations (Minoux, 1978; Leskovec et al., 2007), leveraging on the *submodularity* of σ . More precisely, in Algorithm 2, instead of taking the arg max in the step

$$S_k \leftarrow S_{k-1} \cup \left\{ \arg \max_{i \in V \setminus S_{k-1}} \frac{\sigma(\{i\} \cup S_{k-1}) - \sigma(S_{k-1})}{c_i} \right\},$$

we maintain an upper bound ρ (initially ∞) on the marginal gain, sorted in decreasing order. In each iteration k , we evaluate the element on top of the list, say i , and updates its upper bound with the marginal gain at S_{k-1} . If after the update the upper bound is greater than the others, submodularity guarantees that i is the element with the largest marginal gain.

Algorithm 2 (and the approximation factor) can't be used directly in the OIM context, since computing the exact spread σ is #P hard (Chen et al., 2010). However, with Monte Carlo (MC) simulations, it can efficiently reach an arbitrarily close ratio of $\alpha = 1 - 1/e - \varepsilon$, with a high probability $1 - 1/(t \log^2(t))$ (Kempe et al., 2003).

3.4. An alternative to Lazy Greedy: Ratio maximization from sketches

In the previous approach, MC can still be computationally costly, since the marginal contribution have to be re-evaluated each time, using directed reachability computation in each MC instance. There exist efficient alternatives to α -approximately maximize the spread with cardinality constraint, such as TIM from Tang et al. (2014) and SKIM from Cohen et al. (2014). Adapting TIM to our ratio maximization context is not straightforward, since it require to know the seed set size in advance, which is not the case in Algorithm 2. SKIM is more promising since it uses the standard GREEDY in a *sketch space*. We provide an adaptation of SKIM for approximately maximize the ratio (see Appendix G). It uses the bottom- k min-hash sketches of Cohen et al. (2014), with a threshold for the length of the sketches that depends on both $k = \lfloor \varepsilon^{-2} \log(1/\delta) \rfloor$ and the cost, where δ is an upper bound on the probability that the relative error is larger than ε . Exactly as Cohen et al. (2014) proved the approximation ratio for SKIM, this approach reaches a factor of $1 - 1/e - \varepsilon$, with high probability. More precisely, at round $t \geq 2$, we can actually choose to have the approximation with $\delta = 1/(t \log^2(t))$, only adding a $O(\log(t))$ factor in the computational complexity of Algorithm 2 (Cohen et al., 2014), thus remaining efficient.⁶

3.5. Regret bound for Algorithm 1

We provide a gap dependent upper bound on the regret of BOIM-CUCB in Theorem 1. For this, we define, for $i \in V$, the gap

$$\Delta_{i,\min} \triangleq \min_{S \subset V, p_i(S; \mathbf{w}^*) > 0, \Delta(S) > 0} \Delta(S).$$

We also define, with d_k being the out-degree of node k ,

$$p_{i,\max} \triangleq \max_{S \subset V, p_i(S; \mathbf{w}^*) > 0} \sum_{k \in V} d_k p_k(S; \mathbf{w}^*).$$

Theorem 1. *If π is the policy described in Algorithm 1, then*

$$R_{B,\varepsilon}(\pi) = \mathcal{O} \left(\log B \left(\sum_{i \in V} |V| \frac{\lambda^* + d_i p_{i,\max} |V|}{\Delta_{i,\min}} \right) \right).$$

⁶Wang and Chen (2017) uses the notion of approximation regret with a certain fixed probability. Here, we rather fix this probability to 1 and allow for a $O(\log(t))$ factor in the running time.

In addition, if true costs are known, then

$$R_{B,\varepsilon}(\pi) = \mathcal{O}\left(\log B \left(\sum_{i \in V} \frac{d_i p_{i,\max} |V|^2}{\Delta_{i,\min}}\right)\right).$$

A proof of Theorem 1 can be found in Appendix B. Notice that the analysis can be easily used for the non budgeted setting. In this case, it reduces to the state-of-the-art analysis of Wang and Chen (2017), except that we slightly simplify and improve the analysis to replace the factor $\max_{S \subset V} \sum_{k \in V} d_k \mathbb{I}\{p_k(S; \mathbf{w}^*) > 0\}$ by a potentially much lower quantity $p_{i,\max}$. In the case this last quantity is still large, we can further improve it by considering slight modifications to the original Algorithm 1. This is the purpose of the next section.

4. More refined optimistic spreads

We observe that the factor $p_{i,\max}$ in Theorem 1 can be as large as $|E|$ in the worst case. In other word, if $\Delta = \min_i \Delta_{i,\min}$, the rate can be as large as

$$\mathcal{O}\left(\log B \left(\frac{\lambda^* |V|^2 + |E|^2 |V|^2}{\Delta}\right)\right).$$

We argue here that we can replace $|E|^2 |V|^2$ by $|E| |V|^3 \log^2(|V|)$. Indeed, leveraging on the mutual independence of random variables W_{ij} , we can hope to get a tighter confidence region for \mathbf{w}^* , and thus a provably tighter regret upper bound (Magureanu et al., 2014; Combes et al., 2015; Degenne and Perchet, 2016). We consider the following confidence region from Degenne and Perchet (2016) (see also Perrault et al. (2019a)) and adapted to our setting.

Fact 1 (Confidence ellipsoid for weights). *For all $t \geq 2$, with probability at least $1 - 1/(t \log^2(t))$,*

$$\sum_{ij \in E} N_{\oplus i, t-1} (w_{ij}^* - \bar{w}_{ij, t-1})^2 \leq \delta(t),$$

where $\delta(t) \triangleq 2 \log(t) + 2(|E| + 2) \log \log(t) + 1$.

For OIM (both budgeted and non budgeted), there is a large potential gain in the analysis using the confidence region given by Fact 1 compared to simply using an Hoeffding based one, like in BOIM-CUCB. More precisely, for classical combinatorial semi bandits, Degenne and Perchet (2016) reduced the gap dependent regret upper bound by a factor $\ell / \log^2(\ell)$, where in our case ℓ can be as large as $|E|$. However, there is also a drawback in practice with such confidence region: computing the optimistic spread might be inefficient, even if an oracle for evaluating the spread is available. Indeed, for a fixed $S \subset V$, the problem of maximizing $\mathbf{w} \mapsto \sigma(S; \mathbf{w})$ over \mathbf{w} belonging to some ellipsoid might be hard, since the objective is not necessarily concave. We can overcome this issue using the following Fact 2 (Wen et al., 2017; Wang and Chen, 2017).

Fact 2 (Smoothness property of the spread). *for all $S \subset V$, and all $\mathbf{w}, \mathbf{w}' \in [0, 1]^E$,*

$$\forall k \in V, |p_k(S; \mathbf{w}) - p_k(S; \mathbf{w}')| \leq \sum_{ij \in E} p_i(S; \mathbf{w}) |w_{ij} - w'_{ij}|.$$

In particular,

$$|\sigma(S; \mathbf{w}) - \sigma(S; \mathbf{w}')| \leq |V| \sum_{ij \in E} p_i(S; \mathbf{w}) |w_{ij} - w'_{ij}|.$$

For $S \subset V$ and $\mathbf{w} \in \mathbb{R}^E$, we define the confidence ‘‘bonus’’ as follows:

$$\text{Bonus}(S; \mathbf{w}) \triangleq |V| \sqrt{\delta(t) \sum_{i \in V, N_{\oplus i, t-1} > 0} d_i \frac{p_i(S; \mathbf{w})^2}{N_{\oplus i, t-1}}}.$$

Notice, we don’t sum on vertices with a zero counter. We compensate this by using the convention $\bar{w}_{ij, t-1} = 1$ when $N_{\oplus i, t-1} = 0$. We can successively use Fact 2, Cauchy-Schwartz inequality, and Fact 1 to get, with probability at least $1 - 1/(t \log^2(t))$,

$$\sigma(S; \mathbf{w}^*) \leq \sigma(S; \bar{\mathbf{w}}_{t-1}) + \text{Bonus}(S; \bar{\mathbf{w}}_{t-1}). \quad (7)$$

In the same way, with probability at least $1 - 1/(t \log^2(t))$, we also have (8).

$$\sigma(S; \mathbf{w}^*) \leq \sigma(S; \bar{\mathbf{w}}_{t-1}) + \text{Bonus}(S; \mathbf{w}^*). \quad (8)$$

Contrary to (7), this ‘‘optimistic spread’’ can’t be used directly by the agent since \mathbf{w}^* is not known.

Although the optimistic spread defined in (7) is now much easier to compute, there is still a major drawback that remains: As a function of $S \subset V$, $\text{Bonus}(S; \bar{\mathbf{w}}_{t-1})$ is not necessarily submodular, so the optimistic spread is itself no longer submodular. This is an issue because submodularity is a crucial property for reaching the approximation ratio $1 - 1/e - \varepsilon$. We propose here several submodular upper bound to Bonus, defined for $S \subset V$ and $\mathbf{w} \in \mathbb{R}^E$:

- Bonus₁ is actually modular, and simply uses the subadditivity (w.r.t. S) of Bonus:

$$\text{Bonus}_1(S; \mathbf{w}) \triangleq |V| \sum_{j \in S} \sqrt{\delta(t) \sum_{i, N_{\oplus i, t-1} > 0} d_i \frac{p_i(\{j\}; \mathbf{w})^2}{N_{\oplus i, t-1}}}.$$

- Bonus₂ uses the subadditivity of the square root:

$$\text{Bonus}_2(S; \mathbf{w}) \triangleq |V| \sum_{i, N_{\oplus i, t-1} > 0} p_i(S; \mathbf{w}) \sqrt{\frac{\delta(t) d_i}{N_{\oplus i, t-1}}}.$$

- Bonus₃ uses $p_i(S; \mathbf{w})^2 \leq p_i(S; \mathbf{w})$, and is submodular as the composition between a non decreasing concave function (the square root) and a monotone submodular function:

$$\text{Bonus}_3(S; \mathbf{w}) \triangleq |V| \sqrt{\delta(t) \sum_{i, N_{\oplus i, t-1} > 0} d_i \frac{p_i(S; \mathbf{w})}{N_{\oplus i, t-1}}}.$$

- Bonus_4 uses Jensen's inequality, and is submodular as the expectation of the square root of a submodular function.

$$\text{Bonus}_4(S; \mathbf{w}) \triangleq \mathbb{E} \left[|V| \sqrt{\sum_{i \in V, N_{\oplus i, t-1} > 0, S^{\mathbf{w}^*}} \frac{\delta(t) d_i}{N_{\oplus i, t-1}}} \right],$$

where $\mathbf{W} \sim \otimes_{ij \in E} \text{Bernoulli}(w_{ij})$.

We can write the following approximation guarantees for the two first bonus:

$$\text{Bonus}(S; \mathbf{w}) \leq \text{Bonus}_1(S; \mathbf{w}) \leq |S| \text{Bonus}(S; \mathbf{w}), \quad (9)$$

$$\text{Bonus}(S; \mathbf{w}) \leq \text{Bonus}_2(S; \mathbf{w}) \leq \sqrt{|V|} \text{Bonus}(S; \mathbf{w}).$$

Notice, another approach to get a submodular bonus is to approximate $p_i(S; \bar{\mathbf{w}}_{t-1})$ by the square root of a modular function (Goemans et al., 2009). However, not only this bonus would be much more computationally costly to build than ours, but also, we would get only a $\sqrt{|V|} \log |V|$ approximation factor, which is worst than the one with our Bonus_2 . Since increasing the bonus by a factor $\alpha \geq 1$ increases the gap dependent regret upper bound by a factor α^2 , we only loose a factor $|V|$ for Bonus_2 , compared to the use of Bonus , which is still better than the CUCB approach. Bonus_1 can also be interesting to use when we have some upper bound guarantee on the cardinality of seed sets used (see subsection 5.1). An approximation factor for Bonus_3 or Bonus_4 doesn't seem interesting, because it would involve the inverse of triggering probabilities. We can, however, further upper bound $\text{Bonus}_3(S; \mathbf{w}^*)$ as follows:

$$\begin{aligned} \text{Bonus}_3(S; \mathbf{w}^*) &= |V| \sqrt{\sum_{i, N_{\oplus i, t-1} > 0} \frac{\delta(t) d_i p_i(S; \mathbf{w}^*)}{N_{\oplus i, t-1}}} \\ &\leq |V| \sqrt{\sum_{j \in S} \sum_{\substack{i \in V, \\ N_{\oplus i, t-1} > 0}} d_i \frac{\delta(t) p_i(\{j\}; \mathbf{w}^*)}{N_{\oplus i, t-1}}} \\ &\leq |V| \sqrt{\delta(t) \sum_{j \in S} |E| \left(\frac{8}{N_{\oplus j, t-1}} \wedge 1 \right)}, \end{aligned}$$

where the last inequality only holds under some high probability event, given by the following Proposition 3, involving counters on the costs and counters on the weights.

Proposition 3. Consider the event defined by $\mathfrak{F}_t \triangleq \{\forall i \in V, N_{\oplus i, t-1} \geq \delta(t)\}$. Then, for all $i, j \in V$,

$$\mathbb{P} \left[\mathfrak{F}_t \text{ and } \frac{\delta(t) p_i(\{j\}; \mathbf{w}^*)}{N_{\oplus i, t-1}} > \frac{8\delta(t)}{N_{\oplus j, t-1}} \right] \leq 1/t^2.$$

We thus define for $S \subset V$,

$$\text{Bonus}_5(S) \triangleq |V| \sqrt{\delta(t) \sum_{j \in S} |E| \left(\frac{8}{N_{\oplus j, t-1}} \wedge 1 \right)}.$$

This bonus is much more convenient since it does not depend anymore on \mathbf{w}^* , and can thus be computed by the agent. Indeed, although the first four bonuses are likely to be tighter than this last submodular Bonus_5 , their dependence in \mathbf{w} forces us to use them for $\mathbf{w} = \bar{\mathbf{w}}_{t-1}$. Even if this doesn't pose any problem in practice, this is more difficult to handle in theory since it would involve optimistic estimates on $p_i(\cdot; \bar{\mathbf{w}}_{t-1})$ itself (see the next section for further details). Actually, we will see that the analysis with Bonus_5 is slightly better than the one we would get with Bonus_2 , since we loose a factor $|V| \log^2(|V|) / \log^2(|E|)$ compared to the use of Bonus . In addition, it allows a much more interesting constant term in the regret upper bound thanks to the suppression of the dependence in \mathbf{w}^* .

Although we can improve the analysis based on Fact 1 and 2, the inequality in this last fact may be less rough in practice (we confirm this in section 7). In that case, we suffer from this roughness, since we actually use Fact 2 to design our bonus. In contrast, BOIM-CUCB only uses it in the analysis, and can therefore adapt to a better smoothness inequality. Thus, we consider BOIM-CUCB₅, where we first compute S_t using the BOIM-CUCB approach, and accept it only if

$$\sigma(S_t; \mathbf{w}_t) \leq \sigma(S_t; \bar{\mathbf{w}}_{t-1}) + \text{Bonus}_5(S_t), \quad (10)$$

otherwise, we chose S_t maximizing $\sigma(S; \bar{\mathbf{w}}_{t-1}) + \text{Bonus}_5(S)$. For technical reason (due to Proposition 3), we replace S_t by $S_t \cup \{j\}$ if it exists a $j \in V$, such that

$$N_{\oplus j, t-1} < \delta(t). \quad (11)$$

We thus both enjoy the theoretical advantages of Bonus_5 and the practical advantages of BOIM-CUCB. We give the following regret bounds for this approach.

Theorem 2. If π is the policy following BOIM-CUCB₅, then

$$R_{B, \varepsilon}(\pi) = \mathcal{O} \left(\log B \left(\sum_{i \in V} |V| \frac{\lambda^* + |V| |E| \log^2 |V|}{\Delta_{i, \min}} + \lambda^* |V|^2 \right) \right).$$

A proof of Theorem 2 can be found in Appendix E. Such analysis also holds in the non-budgeted setting, and maximizing the spread only instead of the ratio, we can build a policy π satisfying the following (with the standard definition of the non-budgeted gaps):

$$R_{T, \varepsilon}(\pi) = \mathcal{O} \left(\log T \sum_{i \in V} \frac{|V|^2 |E| \log^2 |V|}{\Delta_{i, \min}} \right).$$

The regret rate is thus better than the one from Wang and Chen (2017), gaining a factor $|E| / (|V| \log^2(|V|))$.

5. Improvements using Bonus₁ and Bonus₄

In this section, we show that the use of Bonus₁ and Bonus₄ leads to a better regret leading term, at the cost of a large second order term. In the following, we propose BOIM-CUCB₁ (resp. BOIM-CUCB₄), that are the same approach as BOIM-CUCB₅ with Bonus₁(\cdot ; $\bar{\mathbf{w}}_{t-1}$) (resp. Bonus₄(\cdot ; $\bar{\mathbf{w}}_{t-1}$)) instead of Bonus₅, and where condition (11) is replaced by

$$\exists j \in V, N_{\oplus j, t-1} \leq |E|\delta(t).$$

5.1. Bonus₁ for low cardinality seed sets

In many real world scenarios, maximal cardinality of seed set is small compared to $|V|$. Indeed, in the non-budgeted setting, it is limited by m , and it is usually assumed that m is much smaller than $|V|$. In the budgeted setting, we will see in section 6 how to limit the cost of the chosen seeds, and this is likely to also induce a limit on the cardinality of seeds. Using Bonus₁ is more appropriate in this situation, according to the approximation factor (9). We state in Theorem 3 the regret bound for BOIM-CUCB₁.

Theorem 3. *If π is the policy BOIM-CUCB₁, and if all seeds selected have a cardinality bounded by m , then we have*

$$R_{B,\varepsilon}(\pi) = \mathcal{O}\left(\log B \left(\sum_{i \in V} m \frac{\lambda^* + m|V|^2 d_i \log^2(|E|)}{\Delta_{i,\min}} + \lambda^* |V|^2 |E| \right)\right).$$

A proof can be found in Appendix F. As previously, we can state the following non-budgeted version, with seed set cardinality constrained by m :

$$R_{T,\varepsilon}(\pi) = \mathcal{O}\left(\log T \left(\sum_{i \in V} \frac{m^2 |V|^2 d_i \log^2(|E|)}{\Delta_{i,\min}} + |E| |V|^2 \right)\right).$$

Notice, for both settings, there is an improvement in the main term (the gap dependent one), in the case $m \leq \sqrt{|V|}$. However, there is also a higher gap independent term that appears.

5.2. Bonus₄: the same performance as Bonus?

We show here that the regret with Bonus₄ is of the same order as what we would have had with Bonus (which is not submodular). However, Bonus₄ does not have the calculation guarantees of the other bonuses. We state in Theorem 4 the regret bound for the policy BOIM-CUCB₄. Notice that we obtain a bound whose leading term improves by a factor $|E|/\log^2|E|$ that of BOIM-CUCB.

Theorem 4. *If π is the policy BOIM-CUCB₄, then we have*

$$R_{B,\varepsilon}(\pi) = \mathcal{O}\left(\log B \left(\sum_{i \in V} \frac{|V|\lambda^* + |V|^2 d_i \log^2(|E|)}{\Delta_{i,\min}} + \lambda^* |V|^2 |E| \right)\right).$$

The proof is given in Appendix J. As previously, we can state the following non-budgeted version. Notice that the cardinality constrain does not appear in the bound.

$$R_{T,\varepsilon}(\pi) = \mathcal{O}\left(\log T \left(\sum_{i \in V} \frac{|V|^2 d_i \log^2(|E|)}{\Delta_{i,\min}} + |E| |V|^2 \right)\right).$$

In spite of the superiority in terms of regret of the use of Bonus₄, we must point out that, in the worst case, the calculation of this bonus may require a number of sample (and thus a time complexity) polynomial in t , which does not meet the criterion of efficiency that we set ourselves at the beginning of the paper.

6. Knapsack constraint for known costs

In their setting, Wang et al. (2020) considered the relaxed constraint

$$\mathbb{E}[\mathbf{e}_{S \cup \{0\}} \mathbf{c}^*] \leq b, \quad (12)$$

instead of ratio maximization, where the expectation is over the possible randomness of S . When true costs are known to the agent, we can actually combine the two settings: a seed set S can be chosen only if it satisfies (12). In this section, we describe modifications this new setting implies. First of all, the regret definition is impacted, and F_B^* is now maximal for policies respecting the constraint (12) within each round. Naturally, the definitions of λ^* and S^* are also modified accordingly. Otherwise, apart from Algorithm 2, there is conceptually no change in the approaches that have been described in this paper. We now described the modification needed to make Algorithm 2 works in this setting. The same sequence of set S_k is considered, but instead of choosing the set that maximizes the ratio over all $k \in \{0, \dots, |V|\}$, we restrict the maximization to $k \in \{0, \dots, j\}$, where j is the first index such that $\mathbf{e}_{S_j \cup \{0\}}^\top \mathbf{c}^* > b$. If this maximizer is not S_j , then it satisfies the constraint and is output. Else, we output S_j with probability $(b - \mathbf{e}_{S_{j-1} \cup \{0\}}^\top \mathbf{c}^*)/c_j^*$ and S_{j-1} with probability $1 - (b - \mathbf{e}_{S_{j-1} \cup \{0\}}^\top \mathbf{c}^*)/c_j^*$. This way, the expected cost of the output is b . We prove in Appendix K the following Proposition 4, giving an approximation factor of $1 - 1/e$ for the above modification of Algorithm 2.

Proposition 4. *The solution S obtained by the modified Algorithm 2 is such that:*

$$(1 - e^{-1}) \frac{\mathbb{E}[\sigma(S^*)]}{\mathbb{E}[\mathbf{e}_{S^* \cup \{0\}}^\top \mathbf{c}^*]} \leq \frac{\mathbb{E}[\sigma(S)]}{\mathbb{E}[\mathbf{e}_{S \cup \{0\}}^\top \mathbf{c}^*]},$$

where the expectation is over the possible randomness of S, S^* .

7. Experiments




Figure 1. Regret curves with respect to the budget B (expectation computed by averaging over 10 independent simulations).

In this section, we present an experiment for Budgeted OIM. In Figure 1, we plot $\mathbb{E}\left[\sum_{t=1}^{T_B-1} \Delta(S_t)\right]$ with respect to the budget B used, running over up to $T = 10000$ rounds. This quantity is a good approximation to the true regret according to Proposition 1. Plotting the true regret would require to compute F_B^* , which is NP-Hard to do. We consider a subgraph of Facebook network (Leskovec and Krevl, 2014), with $|V| = 333$ and $|E| = 5038$, as in Wen et al. (2017). We take $\mathbf{w}^* \sim U(0, 0.1)^{\otimes E}$ and take deterministic, known costs with $c_0^* = 1$, and $c_i^* = d_i / \max_{j \in V} d_j$. BOIM-CUCB+ is the same approach as BOIM-CUCB₅ with $\text{Bonus}(\cdot; \bar{\mathbf{w}}_{t-1})$ instead of Bonus_5 , ignoring that $\text{Bonus}(\cdot; \bar{\mathbf{w}}_{t-1})$ is not submodular (it is only sub-additive).

We observed that in BOIM-CUCB₁, BOIM-CUCB₄, BOIM-CUCB₅, Condition 10 (with the correct bonus instead of Bonus_5) always holds, meaning that those algorithms coincide with BOIM-CUCB *in practice*, and that the gain only appears through the analysis. We thus plot a single curve for these 4 algorithms in Figure 1. On the other hand, we observe only a slight gain of BOIM-CUCB+ compared to BOIM-CUCB.

Our experiments confirm that Fact 2 is less rough in practice, as we already anticipated. Indeed, our submodular bonuses are not tight enough to compete with BOIM-CUCB, although we gain in the analysis. The slight gain that we have for BOIM-CUCB+ suggests that the issue is not only about the tightness of a submodular upper bound, but rather about the tightness of Fact 2. This is supported by the following observation we made: for the Facebook subnetwork, for 1000

random draws of a seed set and vector pairs in $[0, 0.1]^E$, the ratio of the RHS and the LHS in Fact 2 is each time greater than $0.4|V|$.

In Appendix I, we conducted further experiments on a synthetic graph comparing BOIM-CUCB to BOIM-CUCB-REGULARIZED, which greedily maximizes the regularized spread $S \mapsto \sigma(S; \mathbf{w}_t) - \lambda e_S c_t$, where λ is a parameter to set. We observed that for an appropriate choice of λ , a performance similar to BOIM-CUCB can be obtained.

8. Discussion and Future work

We introduced a new Budgeted OIM problem, taking both the costs of influencers and fixed costs into account in the seed selection, instead of the usual cardinality constraint. This better represents the current challenges in viral marketing, since top influencers tend to be more and more costly. Our fixed cost can also be seen as the time that a round takes: A null fixed cost would mean that reloading the network to get a new independent instance is *free* and *instantaneous*.⁷ Obviously, this is not realistic. We also provided an algorithm for Budgeted OIM under the IC model and the edge level semi-bandit feedback setting.

Interesting future directions of research would be to explore other kinds of feedback or diffusion models for Budgeted OIM. For practical scalability, it would also be good to investigate the incorporation of the linear generalization framework (Wen et al., 2017) into Budgeted OIM. Notice, this extension is *not* straightforward if we want to keep our tighter confidence region. More precisely, we believe that a *linear semi-bandit* approach that is aware of independence between edge observations should be developed (the linear generalization approach of Wen et al. (2017) treats each edge observation as *arbitrary correlated*).

In addition to this, exploring how the use of Fact 2 *in the Algorithm* might be avoided while still using confidence region given by Fact 1 would surely improve the algorithms. One possible way would be to use a Thompson Sampling (TS) approach (Wang and Chen, 2018; Perrault et al., 2020a), where the prior takes into account the mutual independence of weights. However, Wang and Chen (2018) proved in their Theorem 2 that TS gives linear approximation regret for some special approximation algorithms. Thus, we would have to use some specific property of the GREEDY approximation algorithm we use.

⁷In this case, using $|S|$ rounds choosing each time a single different influencer $i \in S$ is better than choosing the whole S in a single round.

Acknowledgements

The research presented was supported by European CHIST-ERA project DELTA, French Ministry of Higher Education and Research, Nord-Pas-de-Calais Regional Council, French National Research Agency project BOLD (ANR19-CE23-0026-04).

References

- Auer, P., Cesa-Bianchi, N., and Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256.
- Bai, W., Iyer, R., Wei, K., and Bilmes, J. (2016). Algorithms for optimizing the ratio of submodular functions. In *International Conference on Machine Learning*, pages 2751–2759.
- Carpentier, A. and Valko, M. (2016). Revealing graph bandits for maximizing local influence. In *International Conference on Artificial Intelligence and Statistics*.
- Chen, W., Wang, C., and Wang, Y. (2010). Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Knowledge Discovery and Data Mining*.
- Chen, W., Wang, Y., and Yuan, Y. (2013). Combinatorial multi-armed bandit: General framework and applications. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 151–159, Atlanta, Georgia, USA. PMLR.
- Chen, W., Wang, Y., and Yuan, Y. (2016). Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *Journal of Machine Learning Research*, 17.
- Cohen, E. (1997). Size-estimation framework with applications to transitive closure and reachability. *Journal of Computer and System Sciences*, 55(3):441–453.
- Cohen, E. (2016). Minhash sketches.
- Cohen, E., Delling, D., Pajor, T., and Werneck, R. F. (2014). Sketch-based influence maximization and computation: Scaling up with guarantees. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 629–638. ACM.
- Cohen, E. and Kaplan, H. (2007). Summarizing data using bottom-k sketches. In *Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing*, pages 225–234. ACM.
- Combes, R., Shahi, M. S. T. M., Proutiere, A., and Others (2015). Combinatorial bandits revisited. In *Advances in Neural Information Processing Systems*, pages 2116–2124.
- Degenne, R. and Perchet, V. (2016). Combinatorial semi-bandit with known covariance. In *Advances in Neural Information Processing Systems 29*, pages 2972–2980. Curran Associates, Inc.
- Ding, W., Qiny, T., Zhang, X.-D., and Liu, T.-Y. (2013). Multi-armed bandit with budget constraint and variable costs.
- Feige, U. (1998). A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)*, 45(4):634–652.
- Fujishige, S. (2005). *Submodular functions and optimization*. Annals of discrete mathematics.
- Goemans, M. X., Harvey, N. J., Iwata, S., and Mirrokni, V. (2009). Approximating submodular functions everywhere. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, pages 535–544. Society for Industrial and Applied Mathematics.
- Gomez-Rodriguez, M., Leskovec, J., and Krause, A. (2012). Inferring networks of diffusion and influence. *ACM Trans. Knowl. Discov. Data*, 5(4):21:1–21:37.
- Goyal, A., Bonchi, F., and Lakshmanan, L. V. (2010). Learning influence probabilities in social networks. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining, WSDM '10*, pages 241–250, New York, NY, USA. ACM.
- Goyal, A., Lu, W., and Lakshmanan, L. V. S. (2011). Simpath: An efficient algorithm for influence maximization under the linear threshold model. In *2011 IEEE 11th International Conference on Data Mining*, pages 211–220.
- Kakade, S. M., Kalai, A. T., and Ligett, K. (2009). Playing games with approximation algorithms. *SIAM Journal on Computing*, 39(3):1088–1106.
- Kempe, D., Kleinberg, J., and Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM.
- Khuller, S., Moss, A., and Naor, J. S. (1999). The budgeted maximum coverage problem. *Information processing letters*, 70(1):39–45.
- Krause, A. and Guestrin, C. (2005). A Note on the Budgeted Maximization of Submodular Functions. Technical Report June, CMU.

- Leskovec, J., Krause, A., Guestrin, C., Faloutsos, C., Faloutsos, C., VanBriesen, J., and Glance, N. (2007). Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '07, pages 420–429, New York, NY, USA. ACM.
- Leskovec, J. and Krevl, A. (2014). SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>.
- Lugosi, G., Neu, G., and Olkhovskaya, J. (2019). Online influence maximization with local observations. In *Proceedings of the 30th International Conference on Algorithmic Learning Theory*, volume 98 of *Proceedings of Machine Learning Research*, pages 557–580, Chicago, Illinois. PMLR.
- Magureanu, S., Combes, R., and Proutiere, A. (2014). Lipschitz bandits: Regret lower bound and optimal algorithms. In *Proceedings of The 27th Conference on Learning Theory*, volume 35 of *Proceedings of Machine Learning Research*, pages 975–999, Barcelona, Spain. PMLR.
- Minoux, M. (1978). Accelerated greedy algorithms for maximizing submodular set functions. *Optimization Techniques*, pages 234–243.
- Mitzenmacher, M. and Upfal, E. (2017). *Probability and computing: randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press.
- Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions-I. *Mathematical Programming*, 14(1):265–294.
- Netrapalli, P. and Sanghavi, S. (2012). Learning the graph of epidemic cascades. *SIGMETRICS Perform. Eval. Rev.*, 40(1):211–222.
- Nguyen, H. and Zheng, R. (2013). On budgeted influence maximization in social networks. *IEEE Journal on Selected Areas in Communications*, 31(6):1084–1094.
- Perrault, P., Boursier, E., Perchet, V., and Valko, M. (2020a). Statistical efficiency of thompson sampling for combinatorial semi-bandits. *arXiv preprint arXiv:2006.06613*.
- Perrault, P., Perchet, V., and Valko, M. (2019a). Exploiting structure of uncertainty for efficient matroid semi-bandits. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5123–5132, Long Beach, California, USA. PMLR.
- Perrault, P., Perchet, V., and Valko, M. (2019b). Finding the bandit in a graph: Sequential search-and-stop. In *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 1668–1677. PMLR.
- Perrault, P., Perchet, V., and Valko, M. (2020b). Covariance-adapting algorithm for semi-bandits with application to sparse outcomes. *Proceedings of Machine Learning Research*, 125:1–33.
- Saito, K., Nakano, R., and Kimura, M. (2008). Prediction of information diffusion probabilities for independent cascade model. In *Knowledge-Based Intelligent Information and Engineering Systems*, pages 67–75, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Streeter, M. and Golovin, D. (2009). An online algorithm for maximizing submodular functions. In *Advances in Neural Information Processing Systems*, pages 1577–1584.
- Tang, Y., Shi, Y., and Xiao, X. (2015). Influence maximization in near-linear time: A martingale approach. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, SIGMOD '15, pages 1539–1554, New York, NY, USA. ACM.
- Tang, Y., Xiao, X., and Shi, Y. (2014). Influence maximization: Near-optimal time complexity meets practical efficiency. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pages 75–86. ACM.
- Vaswani, S., Kveton, B., Wen, Z., Ghavamzadeh, M., Lakshmanan, L. V., and Schmidt, M. (2017). Model-independent online learning for influence maximization. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3530–3539. JMLR.org.
- Vaswani, S., Lakshmanan, L. V. S., and Mark Schmidt (2015). Influence maximization with bandits. In *NIPS workshop on Networks in the Social and Information Sciences 2015*.
- Wang, Q. and Chen, W. (2017). Improving regret bounds for combinatorial semi-bandits with probabilistically triggered arms and its applications. In *Neural Information Processing Systems*.
- Wang, S. and Chen, W. (2018). Thompson Sampling for Combinatorial Semi-Bandits.
- Wang, S., Yang, S., Xu, Z., and Truong, V.-A. (2020). Fast thompson sampling algorithm with cumulative oversampling: Application to budgeted influence maximization. *arXiv preprint arXiv:2004.11963*.

Wen, Z., Kveton, B., Valko, M., and Vaswani, S. (2017). On-line influence maximization under independent cascade model with semi-bandit feedback. In *Neural Information Processing Systems*.

Xia, Y., Qin, T., Ma, W., Yu, N., and Liu, T.-Y. (2016). Budgeted multi-armed bandits with multiple plays. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 2210–2216. AAAI Press.